European Journal of Operational Research 218 (2012) 789–800

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Supply chain design considering economies of scale and transport frequencies

Kerstin Baumgartner [a,1], André Fuetterer [b,2], Ulrich W. Thonemann [b,*]

[a] McKinsey & Company, Inc., Herrengasse 1-3, 1010 Vienna, Austria

[b] Department of Supply Chain Management & Management Science, University of Cologne, Albertus-Magnus-Platz, 50923 Cologne, Germany

ARTICLE INFO

Article history:
Received 3 July 2009
Accepted 22 November 2011
Available online 1 December 2011

Keywords:
Supply chain design
Economies of scale
Transport frequencies
Iterative linearization

ABSTRACT

In this paper we consider a 3-echelon, multi-product supply chain design model with economies of scale in transport and warehousing that explicitly takes transport frequencies into consideration. Our model simultaneously optimizes locations and sizes of tank farms, material flows, and transport frequencies within the network. We consider all relevant costs: product cost, transport cost, tank rental cost, tank throughput cost, and inventory cost. The problem is based on a real-life example from a chemical company. We show that considering economies of scale and transport frequencies in the design stage is crucial and failing to do so can lead to substantially higher costs than optimal. We solve a wide variety of problems with branch-and-bound and with the efficient solution heuristics based on iterative linearization techniques we develop. We show that the heuristics are superior to the standard branch-and-bound technique for large problems like the one of the chemical company that motivated our research.

1. Introduction

Increasing globalization and vertical disintegration have resulted in a large increase of transportation volumes. One of the industries that has experienced radical changes is the chemical industry. During the last decade, many companies in this industry have established global production sites and use these sites to meet global demands. Bayer (2007), for instance, planned to invest 1.8 billion USD into a Toluylen–Diisocyanat (TDI) production site in Shanghai by 2009, doubling the production capacity of this site to 300,000 tons per year. Similar investments have been made by many other companies. BASF (2007) expanded the capacities of its plants in Nanjing, Ludwigshafen, Antwerp, and Pasir Gudang from 2002 to 2006 and opened a new plant in Pudong in 2006. SABIC (2006, 2007), the Saudi Basic Industries Corporation, a leading manufacturer of chemicals, fertilizers, plastics, and metals, expanded its global network by taking over the petrochemicals business of the Dutch group DSM (2002), Huntsman's UK petrochemical operations (2006), and GE Plastics' US-based operations (2006). In 2006, SABIC used nearly 500 vessels to ship over 8.6 million metric tons of chemicals and gases to more than 90 ports in over 35 countries around the world.

To match global demand with global supply, chemical companies use global supply chains. Typically, products are produced at a few plants and are shipped to regional tank farms, where they are stored. Customer demand is then met utilizing these regional tank farms. In most situations, the tanks are owned by a tank farm operator that rents the tanks to chemical companies. When designing a supply chain, companies must make a number of decisions: they must decide on the product mix and production quantities at the production sites, on the locations of the tank farms and tank capacities, and on the frequency of deliveries between plants. These supply chain design decisions are medium- to long-term decisions and are typically made annually or bi-annually.

The cost structure of chemical supply chains exhibits several important economies of scale that have to be considered when designing the supply chain. The most important ones are economies of scale in transportation quantities and economies of scale in tank capacities. The freight rate between Europe and South America, for instance, decreases from about 400 USD/m^3 to 200 USD/m^3 when the transportation volume increases from 1000 m^3 to 10,000 m^3. The costs of tank rentals exhibit similar economies of scale. The rental cost of a typical 500 m^3 tank, for instance, is 54 USD/m^3/year, whereas the cost of a 2000 m^3 tank is only 24 USD/m^3/year.

Previous research has addressed some of the issues that are relevant when designing supply chains like the one we consider. However, to our best knowledge, there does not exist an approach that considers all of the main characteristics simultaneously: design of product mixes, transportation routes, transportation frequencies, storage locations, and storage sizes, taking into account non-linearities in transportation and storage costs. Since the decisions greatly impact one another, they must be made simultaneously to obtain an efficient supply chain design and cannot be

* Corresponding author. Tel.: +49 221 470 7935; fax: +49 221 470 7950.
E-mail addresses: kerstin_baumgartner@mckinsey.com (K. Baumgartner), andre.fuetterer@uni-koeln.de (A. Fuetterer), ulrich.thonemann@uni-koeln.de (U.W. Thonemann).
[1] Tel.: +43 1 5370 6156; fax: +43 1 5370 6157.
[2] Tel.: +49 221 470 7938; fax: +49 221 470 7950.

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.11.032

List of symbols

I	set of plants
i	plant
P	set of products
p	product
J	set of tank farm locations
j	tank farm location
K	set of customers
k	customer
t	transport schedule
$q_{ijpt}^{(1)}$	volume of product p shipped from plant i to tank farm j on schedule t
C_{ip}	unit production cost of product p at plant i
Z^P	total production cost
$O_t^{(1)}$	shipping frequency for transport from plants to tank farm locations on schedule t
$\widetilde{Z}_{ijt}^{T1}$	transport cost per trip between plant i and tank farm j on schedule t
N	set of cost levels/cost-function segments in $\widetilde{Z}_{ijt}^{T1}$
n	cost level/cost-function segment in $\widetilde{Z}_{ijt}^{T1}$
$B_{ijn}^{(1)}$	quantity at which the nth breakpoint in $\widetilde{Z}_{ijt}^{T1}$ occurs
$V_{ijn}^{(1)}$	slope/variable cost of the nth segment in $\widetilde{Z}_{ijt}^{T1}$
$F_{ijn}^{(1)}$	Y-intercept/fixed-cost equivalent of the nth segment in $\widetilde{Z}_{ijt}^{T1}$
Z^{T1}	total annual cost for transport between plants and tank locations
$q_{jp}^{(3)}$	volume of product p stored at tank farm location j
$\widetilde{Z}_{jp}^{S1}$	tank rental cost for product p at tank location j
M	set of cost levels/cost-function segments in $\widetilde{Z}_{jp}^{S1}$
m	cost level/cost-function segment in $\widetilde{Z}_{jp}^{S1}$
$B_{jm}^{(3)}$	quantity at which the nth breakpoint in $\widetilde{Z}_{jp}^{S1}$ occurs
$V_{jm}^{(3)}$	slope/variable cost of the nth segment in $\widetilde{Z}_{jp}^{S1}$
$F_{jm}^{(3)}$	Y-intercept/fixed-cost equivalent of the nth segment in $\widetilde{Z}_{jp}^{S1}$
Z^{S1}	Total tank rental cost across all tank locations and products
R_j	tank turn cost for tank location j per volume
Z^{S2}	total tank turn cost across all tank locations and products
H	annual inventory holding cost rate
Z^{S3}	total inventory holding cost across all tank locations and products
Z^S	total storage cost
$O^{(2)}$	shipping frequency for transport from tank farm locations to customers
$q_{jkp}^{(2)}$	volume of product p shipped from tank farm j to customer k
$\widetilde{Z}_{jk}^{T2}$	transport cost per trip between tank farm j and customer k
$V_{jk}^{(2)}$	slope/variable cost in $\widetilde{Z}_{jk}^{T2}$
$F_{jk}^{(2)}$	Y-intercept/fixed-cost in $\widetilde{Z}_{jk}^{T2}$
Z^{T2}	total annual cost for transport between tank locations and customers
$Z(q)$	general cost function (for explanatory purpose)
q	quantity (explanatory)
B_n	quantity at which the nth breakpoint in $Z(q)$ occurs
V_n	slope/variable cost of the nth segment in $Z(q)$
F_n	Y-intercept/fixed-cost equivalent of the nth segment in $Z(q)$
q_n	quantity in the nth segment of $Z(q)$
x_n	binary decision variable, 1 if $q_n > 0$, otherwise 0
D_{kp}	demand for product p from customer k
$\widehat{Z}^{T1}$	linearized version of $\widetilde{Z}^{T1}$
$\widehat{Z}^{S1}$	linearized version of $\widetilde{Z}^{S1}$
$\widehat{Z}^{S}$	linearized version of Z^S
$\widehat{Z}^{T2}$	linearized version of $\widetilde{Z}^{T2}$
$W_{ijt}^{(1)}$	mixed-cost coefficient in $\widehat{Z}^{T1}$
$W_{jk}^{(2)}$	mixed-cost coefficient in $\widehat{Z}^{T2}$
$W_{jp}^{(3)}$	mixed-cost coefficient in $\widehat{Z}^{S1}$
$h_{ij}^{(1)}$	total flow of all products from tank farm j to customer k

made independently. The main contribution of the paper is that it provides an approach for modeling and solving such complex supply chain design problems with economies of scale and transport frequencies and that it provides numerical results from a real-world application.

The resulting model is a non-convex piecewise linear network flow problem, which is known to be NP-hard (e.g. Kim and Pardalos, 2000b). We apply the technique of Balakrishnan and Graves (1989) to state our model as a mixed-integer program. To solve problems of realistic size efficiently, we developed new heuristic solution methods. Our work is based on a project for a chemical company that redesigned its sea freight supply chain and we base our numerical analyses on data of this company. However, the application of the model and the solution approaches are not limited to this setting. They can also be used for designing supply chains in other industries that exhibit similar characteristics as the chemical industry, such as the coal, metal, stone, oil, and gas industries.

The remainder of the paper is organized as follows. In Section 2, we review the relevant literature. In Section 3, we model the supply chain design problem mathematically. We place particular emphasis on modeling non-linearities in costs accurately and on incorporating delivery frequencies as decision variables. In Section 4, we develop several heuristics for solving real-world supply chain design problems. We also show how the optimal solutions can be computed for small to medium-sized problems and how a lower bound on the performance of the optimal policy can be computed. In Section 5, we solve a number of problems and compare the performances of the heuristics. Our results indicate that the heuristics generate solutions that are close to optimality. We also show the advantages of our comprehensive model versus simplified models. In Section 6, we draw conclusions.

2. Literature review

Related to our research are the literature on supply chain design models with non-linear storage cost and linear transportation cost, models with linear storage cost and non-linear transportation cost, and models with non-linear storage and non-linear transportation cost. We review the corresponding literature in this section.

Models with *non-linear storage cost and linear transportation cost* have been developed for uncapacitated and capacitated sites. Feldman et al. (1966) introduced an uncapacitated model that they solved with add and drop heuristics. Refinements of the add and drop heuristic for uncapacitated problems have been developed by Spielberg (1969a,b), Drysdale and Sandiford (1969), Khumawala and Kelly (1974), and Whitaker (1985). Baxter (1984) used a continuous version of the problem, where warehouses can be located anywhere and are not restricted to a set of predefined potential locations. He solved the problem with an adaptive location-allocation and perturbation method. Optimal solution approaches have been developed by Efroymson and Ray (1966), who used branch-and-bound to solve the problem, and by Broek et al. (2006), who relied on Lagrangian relaxation.

One of the first capacitated models was introduced by Kelly and Khumawala (1982). They modeled the problem as a transportation problem and solved this problem iteratively using marginal costs that depend on the current volume allocated to each site. Iterative linearization techniques have been used by Klincewicz (1985) for a multi-product problem and by Klincewicz et al. (1988) for a multi-period problem. Holmberg (1984, 1994) and Holmberg and Ling (1997) developed various solution techniques for capacitated problems with staircase warehousing cost functions. Correia and Captivo (2003) solved a modular capacitated location problem using Lagrangian relaxation. A stochastic model that also includes inventory cost was analyzed by Romeijin et al. (2007).

Models with *linear storage cost and non-linear transportation cost* treat warehousing cost quite differently. In a number of models, warehousing cost is excluded. Klincewicz (1990) analyzed the design of a multi-product, 3-echelon distribution system and developed an exact solution for several special cases. Lapierre et al. (2004) considered a location model with comprehensive cost functions, for instance taking into account full truckload discounts and various transportation modes. They solved the model with a combination of tabu and variable neighborhood search algorithms. O'Kelley and Bryan (1998) analyzed a single product, multi-echelon hub location problem with non-linear transportation cost between hubs and solved a small problem with a commercial solver. Klincewicz (2002) applied a variety of solution approaches, such as enumeration, tabu search, and greedy random adaptive search to solve a hub location problem with hubs on two levels, encouraging interhub flows through discounts. Kim and Pardalos (2000a) developed a slope-scaling and domain contraction approach for solving hub location problems with continuous piece-wise linear cost functions. Kim and Pardalos (2000b) extended the approach to discontinuous functions. A model that includes fixed but no variable warehousing cost has been analyzed by Lin et al. (2006). They determined transportation cost based on total arc flow and assumed that it was independent of the order size. Models that add linear warehousing cost in addition to a fixed charge have been developed by Gümüs and Bookbinder (2004) for a cross-docking model that combines less-than-truckload shipments into full truckloads and by Syam (2002) for capacitated storage locations. Syam (2002) also optimized transport frequencies.

Models with *non-linear storage cost and non-linear transportation cost* were first analyzed by Baumol and Wolfe (1958). They analyzed a model with non-linear storage and transportation cost and showed how locally optimal solutions can be determined. Their model did not include inventory cost or replenishment frequencies. Similar models were solved by Zangwill (1968), who proposed a minimum cost network flow algorithm, and by Soland (1974), who used a branch-and-bound algorithm for solving the problem. Paraschis (1989) included inventory cost in his model but assumed that replenishment frequencies are exogenously given. Fleischmann (1993) expanded this model to a stochastic environment with safety stock. Both, Paraschis (1989) and Fleischmann (1993), used flow models and solved the problem heuristically.

The literature review shows that existing research has addressed subsets of the factors that are relevant when designing supply chains for situations like the one faced by the company that motivated our research. However, previous research has not addressed all of these simultaneously, i.e., optimization of product mix, locations, material flows, and transport frequencies in an environment with economies of scale in transport and warehousing. We fill this gap in this paper.

3. Model formulation

Subsequently, we will state the model of the supply chain we consider: let I be the set of all plants indexed in i and P be the set of all products indexed in p. Each plant in I can produce a subset of products from P, which can be shipped via a set of J potential tank farm locations indexed in j. From there, they can be shipped to serve the demand of a set K of customers indexed in k. (Fig. 1).

In the application that motivated our research, the tank farms are supplied by plants via sea transport, where the transportation cost between plants and tank farms is highly non-linear. The products are then stored in the tanks of the tank farms. The unit rental cost of large tanks is much less than the unit rental cost of small tanks, i.e., the storage cost is highly non-linear. Customers are supplied from tank farms via regional sea freight operators that experience economies of scale in terms of fixed-cost degression, albeit with constant cost factors. We will use superscripts (1), (2), and (3) to denote parameters and variables associated with transportation from plants to tank farm locations, transportation from tank farms to customers, and storage at the tanks, respectively.

3.1. Transportation cost between production sites and tank farms

Products are shipped from production sites to tank farms, typically via liquid bulk ships. Since there exist large economies of scale in transportation, transportation cost exhibits a highly non-linear structure. A typical example is the transportation cost from Rotterdam to Buenos Aires. The logistics service provider that is currently used by the chemical company charges a fixed price of 454,300 USD for volumes of up to 1100 m^3 and 413 USD/m^3 for volumes between 1100 m^3 and 2000 m^3. Volumes above 2000 m^3 and below 5000 m^3 receive an additional unit discount and are priced at 227 USD/m^3. Volumes above 5000 m^3 receive a further additional unit discount and are charged at 97 USD/m^3. Note that the discount only applies to the incremental volume, e.g., for a shipping volume of 2500 m^3, 1100 m^3 are charged at a bulk rate of 454,300 USD, 900 m^3 are charged at 413 USD/m^3 and 500 m^3 are charged at 227 USD/m^3. In practice, this type of discount is usually referred to as *incremental discount*. For this example, Fig. 2 shows how the transportation cost depends on the transportation volume.

The volume per shipping depends in the chosen shipping frequency. Therefore, it is important to include frequencies already in the strategic planning of the supply chain in order to correctly evaluate economies of scale and required tank sizes. Failing to incorporate these factors can lead to suboptimal supply chain designs, as we will show in Section 5.2. In our model, we will therefore use the notion of a set T of transport schedules indexed in t. The transport schedule is an index that specifies the frequency at which a tank farm is served from a plant. In theory, more than one schedule could be used for a product along a certain route. However, in practice we find that it is generally optimal to choose one schedule in order to exploit economies of scale. Table 1 shows the schedules that the chemical company has been considering and reports the corresponding frequencies $O_t^{(1)}$ of the shipments from plants to tank farm locations. In the optimization process, different frequencies can be chosen for different products, customers, and plants.

We denote the volume of product p that are shipped from production site i to tank farm j at a schedule t as $q_{ijpt}^{(1)}$. Since different products are usually shipped in separate tanks but on a single ves-

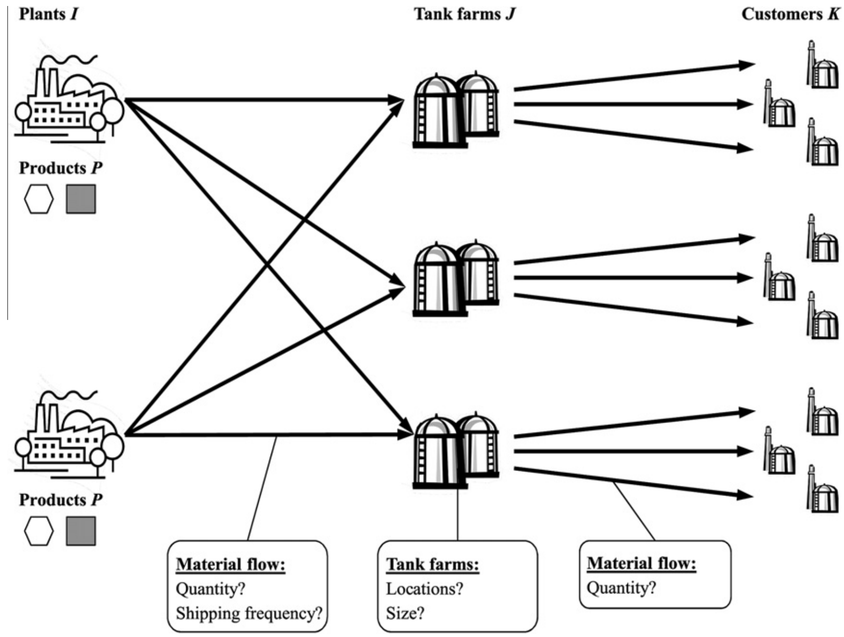

Fig. 1. Schematic supply chain structure.

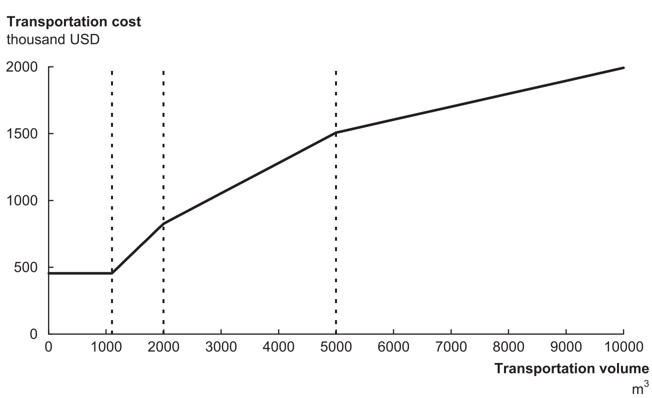

Fig. 2. Transportation cost between Rotterdam and Buenos Aires.

Table 1
Delivery schedules.

t	1	2	3	4
Frequency $O_t^{(1)}$	Bi-weekly 26/yr	Monthly 12/yr	Bi-monthly 6/yr	Quarterly 4/yr

sel, the freight rate is charged based on the total volume of all products shipped per trip, i.e., $\sum_p q_{ijpt}^{(1)}/O_t^{(1)}$.

We model the transportation cost $\widetilde{Z}_{ijt}^{T1}(q_{ijpt}^{(1)})$ per trip between production sites i and tank farms j on transport schedules t as a piecewise linear function. We denote quantities at which breakpoints in the cost function for transport from production site i to tank farm j occur with $B_{ijn}^{(1)}$, where $n \in \{1,\ldots,N\}$ refers to the cost level (i.e., the segment of the piecewise linear function) and consecutively numbers the breakpoints. Since we consider an uncapacitated model, we set $B_{ijN}^{(1)}$ to a "Big M." Let $V_{ijn}^{(1)}$ and $F_{ijn}^{(1)}$ denote the slope and the intercept of the nth segment of this function. In our setting, $V_{ijn}^{(1)}$ are the variable transportation costs per m^3 from plant i to tank farm j at cost level n and $F_{ijn}^{(1)}$ can be directly derived from the corresponding fixed cost. Note that shipping companies charge each booked schedule separately, meaning that when two schedules happen to coincide, there is no extra discount given due to the larger volume on this particular shipping. The two bookings are charged as separate shippings. The transportation cost per trip from a production site i to a tank farm j for a specific schedule t is

$$\widetilde{Z}_{ijt}^{T1}(q_{ijpt}^{(1)}) = \begin{cases} F_{ijn}^{(1)} + V_{ijn}^{(1)} \frac{\sum_p q_{ijpt}^{(1)}}{O_t^{(1)}} & \text{if } B_{ij,n-1}^{(1)} < \frac{\sum_p q_{ijpt}^{(1)}}{O_t^{(1)}} \leqslant B_{ijn}^{(1)} \\ 0 & \text{if } \sum_p q_{ijpt}^{(1)} = 0 \end{cases} \quad \forall i,j,t. \tag{1}$$

The total annual transportation cost for all production sites i, tank farms j, and transport schedules t can be calculated as

$$Z^{T1}(q_{ijpt}^{(1)}) = \sum_{i,j,t} \left(\widetilde{Z}_{ijt}^{T1}(q_{ijpt}^{(1)}) O_t^{(1)} \right). \tag{2}$$

We show how to transform Eq. (1) into mixed-integer form in Section 3.5.

3.2. Production cost

The volume of product p that is produced at production site i can be computed as $\sum_{j,t} q_{ijpt}^{(1)}$ – which is the total volume of product p shipped from a plant i to all tank farms j for all schedules t or, in other words, the total outbound shippings of the respective plant. Note that in general only one t will be active on a specific route for a specific product so that economies of scale can be exploited, although we do not include constraints to hinder more than one transport schedule from being chosen. We sum over all t because ships headed to different locations j can have different schedules. Since different production sites have different labor costs, efficiencies, equipment, etc., the unit production cost differs according to production site. Let C_{ip} denote the unit production cost of product p at site i. Then the total production cost Z^P can be computed as

$$Z^P(q_{ijpt}^{(1)}) = \sum_{i,p} \left(C_{ip} \sum_{j,t} q_{ijpt}^{(1)} \right). \tag{3}$$

3.3. Storage cost

Liquids are shipped from production sites to tank farms, where they are stored before distribution to the customers. A separate tank is needed for each product to prevent liquids from mixing. Three types of costs are incurred for storage: the cost for renting a tank, the cost for tank turns, and inventory holding cost.

Tank rental cost is subject to economies of scale, since tank farm owners charge lower fees per m^3 for larger tanks than for smaller tanks. The tank rental cost structure of the tank farm located in the harbor of Durban in South Africa is a typical example where unit tank rental cost is reduced by 50% if sufficiently large volumes are rented. Table 2 shows the detailed rate structure.

We denote the amount of product p stored at tank farm location j and thus the required tank size with $q^{(3)}_{jp}$. Note that these are auxiliary variables, since the stored amount is equal to the size of all shipments arriving at a tank within a shipping cycle, i.e., $\sum_{i,t}\left(q^{(1)}_{ijpt}/O^{(1)}_t\right)$. $B^{(3)}_{jm}$ denotes the storage quantities at which breakpoints in the tank rental cost function of tank farm j occur, where $m \in \{1,\ldots,M\}$ refers to the respective cost level and consecutively numbers the breakpoints. Since we use an uncapacitated model, we set $B^{(3)}_{jM}$ to a "Big M." Tank rental cost at location j consists of a fixed annual fee incurred per tank rented, $F^{(3)}_{jm}$, and of a variable cost portion charged annually per m^3 of the rented tank, $V^{(3)}_{jm}$. The resulting tank rental cost for a location j and product p is given by

$$\widetilde{Z}^{S1}_{jp}(q^{(3)}_{jp}) = \begin{cases} F^{(3)}_{jm} + V^{(3)}_{jm} q^{(3)}_{jp} & \text{if } B^{(3)}_{j,m-1} < q^{(3)}_{jp} \leqslant B^{(3)}_{jm}, \\ 0 & \text{if } q^{(3)}_{jp} = 0, \end{cases} \quad \forall j,p. \tag{4}$$

In total, the tank rental cost across all tank locations j and products p is calculated as

$$Z^{S1} = \sum_{j,p} \widetilde{Z}^{S1}_{jp}(q^{(3)}_{jp}). \tag{5}$$

We show how to transform Eq. (4) into mixed-integer form in Section 3.5.

Tank turn cost is charged on top of the rental fees. One tank turn is included in the rental fee. For additional tank movements at location j (i.e., inflows and outflows), a variable fee R_j based on volume is charged. The relevant quantity per location j and product p is calculated by subtracting the included tank turn (i.e., $q^{(3)}_{jp}$, the size of the tank) from total throughput, i.e., $\sum_{i,t} q^{(1)}_{ijpt}$. The resulting cost for all locations j and products p is

$$Z^{S2}(q^{(1)}_{ijpt}, q^{(3)}_{jp}) = \sum_{j,p}\left(\left(\sum_{i,t} q^{(1)}_{ijpt} - q^{(3)}_{jp}\right) R_j\right). \tag{6}$$

Inventory holding cost is calculated based on average working capital. The average volume of stored product p coming from production site i is $\sum_{j,t}\left(q^{(1)}_{ijpt}/(2O^{(1)}_t)\right)$. To calculate the holding cost, we need to obtain the monetary value of the average inventory level based on the product cost C_{ip} and charge an annual holding fee H, which is given by the company's cost of capital:

Table 2
Example warehouse rental rates.

Quantity	Fixed rental rate	Variable rental rate
Less than 900 m^3	4800 USD	44 USD/m^3
Between 900 m^3 and 1350 m^3	4800 USD	31 USD/m^3
Greater than 1350 m^3	4800 USD	22 USD/m^3

$$Z^{S3}(q^{(1)}_{ijpt}) = \sum_{i,p}\left(\sum_{j,t}\left(q^{(1)}_{ijpt}/(2O^{(1)}_t)\right)C_{ip}\right)H. \tag{7}$$

By adding the three components of the storage cost, we obtain the total storage cost

$$Z^{S}(q^{(1)}_{ijpt}, q^{(3)}_{jp}) = Z^{S1}(q^{(3)}_{jp}) + Z^{S2}(q^{(1)}_{ijpt}, q^{(3)}_{jp}) + Z^{S3}(q^{(1)}_{ijpt}). \tag{8}$$

3.4. Transportation cost between tank farms and customers

Customers are supplied from tank farms on a fixed shipping schedule. These shipping schedules are given by customer demands. They are more frequent and volumes are smaller in comparison to the shipments from the plants to the tank farms. We use the parameter $O^{(2)}$ to denote the number of shippings per year, i.e., we assume that all customers are served with the same schedule. This assumption holds for the chemical company that motivated our research, where customers are served with a weekly schedule. We comment in Section 6 on how the model would change and how it could be solved in situations where this assumption does not hold. We denote the quantity of product p shipped from tank farm j to customer k with $q^{(2)}_{jkp}$. Different fixed and variable transportation costs apply on each route from tank farm j to customer k and we denote them with $F^{(2)}_{jk}$ and $V^{(2)}_{jk}$, respectively. Note that there are no volume discounts in the tariff structure. Economies of scale only exist in the form of fixed-cost degression. We discuss how volume discounts for this transport step could be incorporated in Section 6. Since multiple products p can be shipped on one vessel by using separate tanks, the volume of each shipment going from tank farm j to customer k can be computed as $\sum_p q^{(2)}_{jkp} \big/ O^{(2)}$. The resulting transportation cost from tank farm j to customer k is

$$\widetilde{Z}^{T2}_{jk}(q^{(2)}_{jkp}) = \begin{cases} F^{(2)}_{jk} + V^{(2)}_{jk}\left(\sum_p q^{(2)}_{jkp} \Big/ O^{(2)}\right) & \text{if } \sum_p q^{(2)}_{jkp} > 0, \\ 0 & \text{otherwise}, \end{cases} \quad \forall j,k. \tag{9}$$

The total transport cost from all tank farms to all customers can be computed as

$$Z^{T2}(q^{(2)}_{jkp}) = \sum_{j,k} \widetilde{Z}^{T2}_{jk}(q^{(2)}_{jkp}) O^{(2)}. \tag{10}$$

In the next section, we will transform this notation into mixed-integer form, interpreting the function as a piecewise linear cost func-

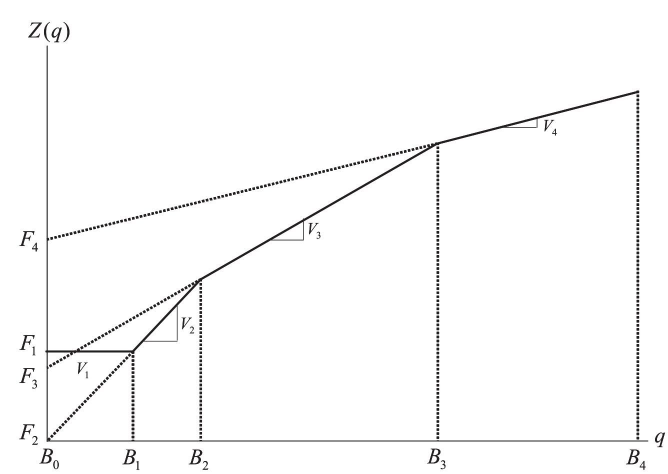

Fig. 3. Example for modeling technique of piecewise linear functions.

tion with one piece. Note that we include a charge for a fixed fee for using a route. This cost structure models well the situation of the chemical company, which uses a third-party shipping company for serving its customers under a tariff with a fixed fee for each delivery.

3.5. Linearization of transportation and storage cost

Parts of the supply chain costs stated before are characterized by piecewise linear cost functions. We now introduce the general method we use to model these piecewise linear cost functions in mixed-integer form. This methodology was first proposed by Balakrishnan and Graves (1989). In general, let $Z:\mathbb{R}\to\mathbb{R}$ be an arbitrary, piecewise linear cost function of quantity q that consists of N different, linear sections (i.e., cost levels), defined by

$$Z(q)=\begin{cases}F_n+V_nq & \text{if } B_{n-1}<q\leqslant B_n,\\ 0 & \text{if } q=0,\end{cases} \tag{11}$$

where $0=B_0<B_1<\cdots<B_n<\cdots<B_N$ denote the breakpoints between the nth and $(n+1)$th segment and F_n and V_n denote the intercept and the slope of the nth segment. Fig. 3 gives an example of such a function.

In order to transform Eq. (11) into mixed-integer form, we introduce q_n as the quantity in segment n, i.e.,

$$q_n=\begin{cases}q & \text{if } B_{n-1}<q\leqslant B_n,\\ 0 & \text{otherwise},\end{cases}\quad \forall n \tag{12}$$

We also define a binary variable x_n that, in our setting, represents the decision to open or close some route or warehouse at cost level n, i.e.,

$$x_n=\begin{cases}1 & \text{if } q_n>0,\\ 0 & \text{otherwise},\end{cases}\quad \forall n. \tag{13}$$

The mixed integer model that results for a given quantity q is

$$Z(q)=\sum_n(F_nx_n+V_nq_n), \tag{14}$$

$$\sum_n q_n=q, \tag{15}$$

$$\sum_n x_n\leqslant 1, \tag{16}$$

$$q_n\leqslant x_nB_n\quad \forall n, \tag{17}$$

$$q_n\geqslant x_nB_{n-1}\quad \forall n, \tag{18}$$

$$q_n\geqslant 0\quad \forall n, \tag{19}$$

$$x_n\in\{0,1\}\quad \forall n. \tag{20}$$

Function (14) calculates the cost for a piecewise linear cost structure depending on the volume. Constraint (15) states that the total quantity q needs to be distributed among different cost levels n by assigning portions of this quantity q to different q_n. The cost level essentially specifies which piecewise linear part of the cost curve is used and depends on the volume. Constraint (16) states that at most one x_n may be set equal to 1. For functions that continuously exhibit economies of scale, this constraint is not necessary because it is always more favorable to pool volumes in order to capture economies of scale. Constraint (17) states that the quantity assigned to q_n may not exceed the upper bound B_n of the cost function for the respective cost level n. Similarly, Constraint (18) states that the quantity may not be lower than the lower bound B_{n-1} of the cost function for the respective cost level n. (18) is only necessary for non-concave functions. For concave functions, it is never optimal to select a cost level lower than n for quantities in the range of cost level n. Constraints (17) and (18) also force the x_n that corresponds to the $q_n>0$ to be set to 1. Constraints (19) and (20) are standard non-negativity and integrality constraints. This model formulation for piecewise linear functions automatically assigns the volume to the correct portion of the cost function. We will apply this technique to transform Eqs. (1), (4), and (9) into mixed-integer form. Next, we present the mathematical model of our supply chain design problem.

3.6. Mathematical model

The chemical company was faced with the challenge of finding the right trade-off between the four cost factors introduced above in order to meet the demand D_{kp} of each customer k for all products p at optimal cost. We will now describe the mathematical model used to solve this complex problem.

The objective is to minimize all relevant supply chain costs:

$$\min_{q^{(1)}_{ijpt},q^{(2)}_{jkp},q^{(3)}_{jp}} Z^P(q^{(1)}_{ijpt})+Z^{T1}(q^{(1)}_{ijpt})+Z^S(q^{(1)}_{ijpt},q^{(3)}_{jp})+Z^{T2}(q^{(2)}_{jkp}), \tag{21}$$

subject to (3)–(10), where all transformations of cost functions into mixed-integer form are performed along the logic of (14)–(20), and subject to

$$\sum_j q^{(2)}_{jkp}=D_{kp}\quad \forall k,p, \tag{22}$$

$$\sum_{i,t} q^{(1)}_{ijpt}=\sum_k q^{(2)}_{jkp}\quad \forall j,p, \tag{23}$$

$$q^{(3)}_{jp}=\sum_{i,t}\left(q^{(1)}_{ijpt}/O^{(1)}_t\right)\quad \forall j,p, \tag{24}$$

$$q^{(1)}_{ijpt},q^{(2)}_{jkp},q^{(3)}_{jp}\geqslant 0\quad \forall i,j,k,p,t. \tag{25}$$

The objective function (21) comprises the total supply chain cost. Constraint (22) ensures that the annual demand of each customer is met. It requires the sum of all shipments to the customer to equal the respective customer's demand. Constraint (23) ensures that material inflows and outflows match for each tank farm. For each tank farm, the amount of each product shipped to the tank farm from all production sites must be equal to the amount of the product shipped from this tank farm to all customers (24). Constraint (25) is a standard non-negativity constraint.

4. Solution approach

We implement three types of solution approaches: one exact solution method and two heuristic approaches based on different linearization techniques. For the heuristic approaches, we use a deterministic method and a randomized variant. To make the performance of the different approaches comparable, tests are limited to a run time of one hour.

4.1. Direct solution with a standard branch-and-bound code

The first possibility for solving the problem is using a commercial mixed-integer solver. Next to the best integer solution found in the given time frame of one hour, we also record the lower bound supplied by the solver as a benchmark for all solution methods. Throughout the optimization process, lower bounds are determined by the mixed-integer solver with LP relaxations for the subproblems of the branch-and-bound tree. The lower bound we refer to is the lower bound given by the MIP solver after a run time of one hour, i.e., after pruning the branch-and-bound tree.

4.2. Linearization by integrality relaxation

The first type of linearization technique relaxes all integrality constraints in the mixed-integer formulation corresponding to

Constraint (20) and introduces box constraints of form (26); all other components of the model remain unchanged.

$$0 \leqslant x_n \leqslant 1 \quad \forall n. \tag{26}$$

Different forms of volume reallocations and drop mechanisms are utilized to improve the solution. We implement a deterministic heuristic (*D1*) and a heuristic with a random element (*R1*).

D1 – deterministic relaxed MIP drop heuristic: The relaxed mixed-integer model is solved with the CPLEX LP solver. Due to the relaxation, the LP solver determines solutions that entail opening fractions of tanks in a location. This leads to suboptimal and invalid choices in terms of the part of the piecewise linear cost functions utilized. To restore feasibility and integrality, we calculate the correct costs, making sure that the non-linear cost functions are utilized correctly, with Eqs. (1)–(5), and (8)–(10).

Since fractions of tanks are allowed in the relaxed model, which makes it much cheaper to open tanks, there are too many open tanks in the initial solution. A drop heuristic is employed to improve the solution. Tank farms are sorted in ascending order by throughput. Then, location by location, all open tanks at the respective tank farm are fixed to closed and the modified relaxed problem is solved again. In our experience, closing entire locations instead of single tanks leads to better results and is quicker. Once again, the reallocations to the optimal part of each piecewise linear cost function are performed according to Eqs. (1)–(5), and (8)–(10) thus also restoring integrality. If closing the tank farm has improved the solution, it is kept closed. If there is no improvement, the drop step is undone and we proceed with the next tank farm. Once all tank farm locations have been analyzed in this manner, a new sorted list is generated with the remaining open tank farms and the entire process is repeated until no more improvements can be achieved. Algorithm 1 summarizes the procedure.

Algorithm 1: Deterministic relaxed MIP drop heuristic

```
Relax integrality constraints
Solve LP
Calculate correct cost of original mixed-integer model
  before relaxation
repeat
    Sort open tank locations by throughput volume in
  ascending order
    for all Tank locations in list do
        Close location
        Solve LP
        Calculate correct cost of original mixed-integer
  model before relaxation
        if Current solution < best solution then
            Fix location to be closed
        else
            Re-open tank location
        end if
    end for
until No improvement or runtime > 1 hour
```

R1 – randomized relaxed MIP Drop heuristic: This heuristic works like its deterministic counterpart *D1*, except that the tank farms in the drop heuristic are not sorted in ascending order. Instead, they are sorted randomly: Each tank farm is assigned a probability proportional to its share of the total throughput of all tanks. Then, all tank farms are sorted according to their probability and we compute the cumulative probability. We generate a random number u from the uniform distribution on the interval [0,1] and search for the first tank farm in the list whose cumulative probability is greater or equal to u. This tank farm is added to the final sorted list and its probability is set to 0. The cumulative probabilities are updated and the procedure is repeated until all tank farms have been added to the final list. This probabilistic element reduces the likelihood of getting stuck in local minima. The heuristic is repeated until a run time of one hour is reached.

4.3. Linearization by model reformulation

For the second type of linearization technique, we remove all integer variables from the problem formulation and use mixed-cost coefficients that combine fixed and variable cost parts in one coefficient. The resulting LP relaxations can be solved very quickly. In our approach, we solve successive LP problems, in which we iteratively update these cost coefficients based on the volume allocated to a certain route or tank in the current LP solution. The technique is similar to that found in Kim and Pardalos (2000a,b), who apply it to hub location problems.

For the adapted problem formulation, all binary decision variables corresponding to x_n in Model (14)–(20) are removed from the problem. The piecewise linear cost functions Eqs. (1)–(5), and (8)–(10) are replaced by linear functions (28)–(31) with coefficients $W^{(1)}_{ijt}$, $W^{(2)}_{jk}$, and $W^{(3)}_{jp}$. We use $\widehat{Z}^{T1}(q^{(1)}_{ijpt})$, $\widehat{Z}^{S1}(q^{(3)}_{jp})$, $\widehat{Z}^{S}(q^{(1)}_{ijpt}, q^{(3)}_{jp})$, and $\widehat{Z}^{T2}(q^{(2)}_{jkp})$ to denote the cost components of the linearized model.

The modified objective function reads

$$\min_{q^{(1)}_{ijpt},q^{(2)}_{jkp},q^{(3)}_{jp}} Z^P(q^{(1)}_{ijpt}) + \widehat{Z}^{T1}(q^{(1)}_{ijpt}) + \widehat{Z}^{S}(q^{(1)}_{ijpt}, q^{(3)}_{jp}) + \widehat{Z}^{T2}(q^{(2)}_{jkp}), \tag{27}$$

where

$$\widehat{Z}^{T1}(q^{(1)}_{ijpt}) = \sum_{i,j,p,t} \left(q^{(1)}_{ijpt} W^{(1)}_{ijt}\right), \tag{28}$$

$$\widehat{Z}^{S1}(q^{(3)}_{jp}) = \sum_{j,p} \left(q^{(3)}_{jp} W^{(3)}_{jp}\right), \tag{29}$$

$$\widehat{Z}^{S}(q^{(1)}_{ijpt}, q^{(3)}_{jp}) = \widehat{Z}^{S1}(q^{(3)}_{jp}) + Z^{S2}(q^{(1)}_{ijpt}, q^{(3)}_{jp}) + Z^{S3}(q^{(1)}_{ijpt}), \tag{30}$$

$$\widehat{Z}^{T2}(q^{(2)}_{jkp}) = \sum_{j,k,p} \left(q^{(2)}_{jkp} W^{(2)}_{jk}\right). \tag{31}$$

The cost functions (3), (6), and (7) remain unchanged. We now no longer need to linearize with the help of model (14)–(20), since all cost functions are linear already. Constraints (22)–(25) remain unchanged.

Observe that variable and fixed-cost parameters $V^{(1)}_{ijn}$, $V^{(2)}_{jk}$, $V^{(3)}_{jm}$, $F^{(1)}_{ijn}$, $F^{(2)}_{jk}$, and $F^{(3)}_{jm}$ are no longer used in the model formulation. However, we need them when determining the values of the corresponding linearized parameters $W^{(1)}_{ijt}$, $W^{(2)}_{jk}$, and $W^{(3)}_{jp}$. They contain the variable cost factor plus a portion of the fixed cost resulting from spreading the fixed cost evenly across the current volumes. They are updated after each optimization run based on the volumes in the current LP solution. We include the index t in $W^{(1)}_{ijt}$ because it is necessary to track separate mixed-cost parameters for each transport schedule in this heuristic. Similarly, we use the index p in the variable $W^{(3)}_{jp}$ because different mixed-cost levels for tank rentals apply for different products, depending on the quantity stored of each product.

Based on this method, we implement two heuristics. *D2* is deterministic whereas *R2* has a random element.

D2 – deterministic dynamic slope-scaling drop heuristic: To obtain an initial solution, the mixed-cost coefficients are set to initial values according to Eqs. (32)–(34) and the linear model is solved. We spread the fixed cost across the largest theoretically possible volume for the respective route or tank. This theoretical volume is based on overall demand. For transport from production sites to tank farms, this volume is calculated as total customer demand divided by the number of shippings, $\sum_{kp} D_{kp}/O^{(1)}_t$. For transport from tank farms to customers, the maximum possible volume is the respective customer's demand across all products, $\sum_p D_{kp}$. For storage at the tank farms, the largest possible volume per product is the total demand for the respective product divided by the smallest possible number of shippings, $\sum_k D_{kp}/\min_t(O_t^{(1)})$. Note that the transport cost function from tank farm to customer only has one piece.

$$W_{ijt}^{(1)} = V_{ijN}^{(1)} + \frac{F_{ijN}^{(1)}}{\sum_{kp} D_{kp}/O_t^{(1)}} \quad \forall i,j,t, \tag{32}$$

$$W_{jk}^{(2)} = V_{jk}^{(2)} + \frac{F_{jk}^{(2)}}{\sum_p D_{kp}} \quad \forall j,k, \tag{33}$$

$$W_{jp}^{(3)} = V_{jM}^{(3)} + \frac{F_{jM}^{(3)}}{\sum_k D_{kp}/\min_t(O_t^{(1)})} \quad \forall j,p. \tag{34}$$

After solving the LP, the mixed-cost coefficients are updated depending on the volume allocated to each route and on the size of the rented tanks at each tank farm location according to Eqs. (35)–(37). These equations essentially calculate the total volume allocated to the route or tank and then spread the fixed cost equally across this volume, adding this fixed-cost portion to the variable cost. The result is a "variable" cost factor that actually also contains the relevant fixed-cost portion for the current volume. Note that cost factors remain unchanged for those routes and tanks to which no volume was allocated in the current run, i.e., the cost factor from the current run is also used for the following run.

When determining the cost factor $W_{ijt}^{(1)}$ for transport from plants to tank farms, we distinguish between used and unused schedules on those routes that are currently in use. For schedules currently in use $\left(\sum_p q_{ijpt}^{(1)} > 0\right)$, we use the volume transported on this route per trip, $\sum_p q_{ijpt}^{(1)}/O_t^{(1)}$, to determine the correct cost level. To facilitate switching from one transport schedule to another in the optimization process, we assign all unused schedules on a route $\left(\sum_p q_{ijpt}^{(1)} = 0\right)$ with the theoretical cost for that route that would apply if all the volume on the route across all schedules were allocated to the respective schedule. We use the auxiliary variable $h_{ij}^{(1)} = \sum_{p,t} q_{ijpt}^{(1)}$ to calculate the total material flow of all products p from plants i to tank farms j irrespective of schedule t. To evaluate the cost level, we calculate the theoretical volume per trip, $h_{ij}^{(1)}/O_t^{(1)}$ for each unused schedule. This pooling of volumes facilitates switching to another schedule in the next iteration because alternative, currently unused schedules are also assigned attractive costs. It is important to differentiate between used and unused schedules to ensure that the correct transportation cost is used when calculating the cost for the routes actually in use.

$$W_{ijt}^{(1)} = \begin{cases} V_{ijn}^{(1)} + \frac{F_{ijn}^{(1)}}{\sum_p q_{ijpt}^{(1)}/O_t^{(1)}} & \text{if } \sum\limits_p q_{ijpt}^{(1)} > 0 \text{ and } B_{ij,n-1}^{(1)} < \sum\limits_p q_{ijpt}^{(1)}/O_t^{(1)} \leqslant B_{ijn}^{(1)} \\ V_{ijn}^{(1)} + \frac{F_{ijn}^{(1)}}{h_{ij}^{(1)}/O_t^{(1)}} & \text{if } \sum\limits_p q_{ijpt}^{(1)} = 0 \text{ and } B_{ij,n-1}^{(1)} < h_{ij}^{(1)}/O_t^{(1)} \leqslant B_{ijn}^{(1)} \\ \text{unchanged} & \text{otherwise} \end{cases} \quad \forall i,j,t, \tag{35}$$

To determine the cost factor $W_{jk}^{(2)}$ for transport from tank farms to customers, we first check which routes are in use in the current solution $\left(\sum_p q_{jkp}^{(2)} > 0\right)$. For these routes, we spread the fixed cost across the total volume per shipping, $\sum_p q_{jkp}^{(2)}/O^{(2)}$.

$$W_{jk}^{(2)} = \begin{cases} V_{jk}^{(2)} + \frac{F_{jk}^{(2)}}{\sum_p q_{jkp}^{(2)}/O^{(2)}} & \text{if } \sum\limits_p q_{jkp}^{(2)} > 0 \\ \text{unchanged} & \text{if } \sum\limits_p q_{jkp}^{(2)} = 0 \end{cases} \quad \forall j,k, \tag{36}$$

To calculate the cost factor $W_{jp}^{(3)}$, we determine the correct cost level for the storage volume $q_{jp}^{(3)}$ and spread the corresponding fixed cost across this volume.

$$W_{jp}^{(3)} = \begin{cases} V_{jm}^{(3)} + \frac{F_{jm}^{(3)}}{q_{jp}^{(3)}} & \text{if } B_{j,m-1}^{(3)} < q_{jp}^{(3)} \leqslant B_{jm}^{(3)} \\ \text{unchanged} & \text{if } q_{jp}^{(3)} = 0 \end{cases} \quad \forall j,p. \tag{37}$$

After each iteration, the open tank farms are sorted in ascending order by throughput volume. Location by location, all the tanks at each tank farm are closed and the problem is solved again. The mixed-cost coefficients are updated according to the scheme described above. If there is an improvement, the tank farm is fixed closed and the updated cost coefficients are kept. Otherwise, the drop step is undone, the previous cost coefficients are restored, and we proceed with the next tank farm on the list. Once all tank farms have been checked in this manner, a new sorted list with the remaining open tank farms is generated and the whole process is repeated until no more improvements can be achieved. After the drop mechanism terminates, all variables are freed (i.e., no tank farms are fixed closed anymore) and the model is solved once more. Again, if there is a cost improvement, the solution and the updated cost coefficients are kept. This process is repeated until the solution does not change anymore or until a run time of one hour is exceeded. Algorithm 2 provides an overview of the procedure.

Algorithm 2: Deterministic dynamic slope-scaling drop heuristic

Initialize coefficients, i.e. $W_{ijt}^{(1)}, W_{jk}^{(2)}, W_{jp}^{(3)}$
Solve LP
Update coefficients
repeat
  repeat
    Sort *open* tank locations by throughput volume in ascending order
    for all Tank locations in list **do**
      Close location
      Solve LP
      Update coefficients
      if Current solution < best solution **then**
        Fix location to be closed
      else
        Re-open tank location
        Restore coefficients
      end if
    end for
  until No improvement **or** runtime > 1 hour
  Unfix all tank locations
  Solve LP
  Update coefficients
until No improvement **or** runtime > 1 hour

R2 – randomized dynamic slope-scaling drop heuristic: This approach works like the deterministic heuristic *D2*, but occasionally solutions that increase costs are kept in the drop heuristic in order to leave local optima in pursuit of the global optimum. The approach is inspired by Simulated Annealing (Kirkpatrick et al., 1983).[3]

Let the difference between the current solution Z^{cur} and the best solution so far $Z^{\min}$ be $\Delta = Z^{cur} - Z^{\min}$. The probability that a solution is accepted is

[3] However, our approach does not use a real annealing schedule with decreasing "temperature".

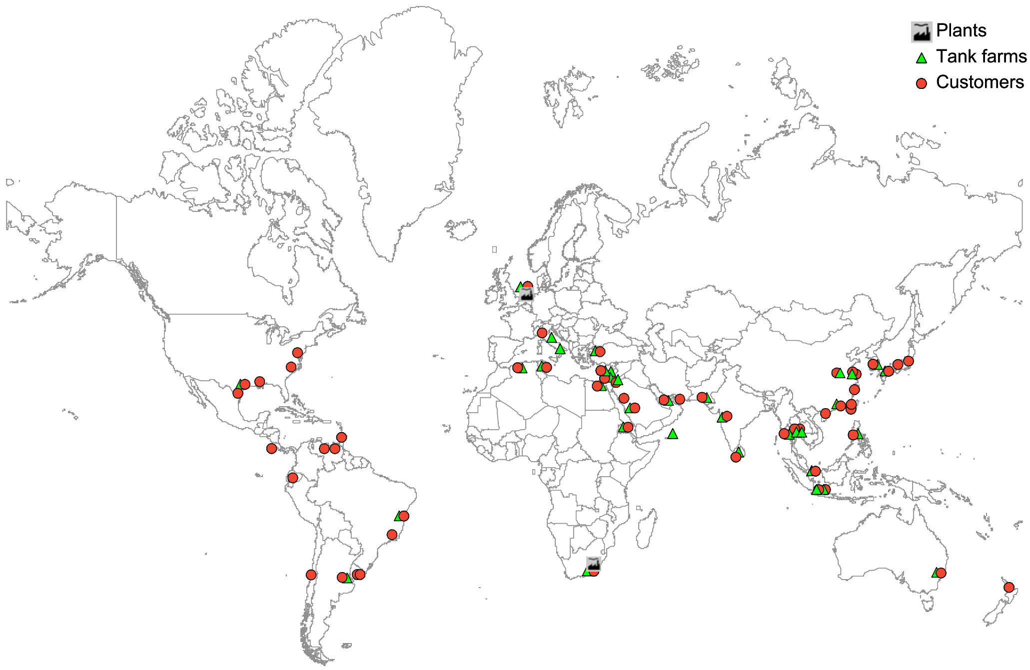

Fig. 4. Overview of locations of real-life supply chain.

$$prob(accept) = \begin{cases} 1 & \text{if } \Delta < 0, \\ e^{-\Delta/g} & \text{if } \Delta \geqslant 0, \end{cases} \quad (38)$$

where g is a control parameter. It determines the probability of non-improving solutions being accepted. We tested different values for g and found 0.06% of the current best solution to be effective. We found that leaving the parameter g constant worked well.

After each drop-step, Δ is calculated and the probability of accepting is determined. The decision on acceptance is made with the help of a generated random number. If a solution is accepted, the mixed cost coefficients associated with it are kept and the tank farm closed in the drop step is fixed closed. The heuristic terminates after a run time of one hour.

5. Computational results

We base all our computations on data from the chemical company. The company operates 2 plants, uses 35 potential tank farms, serves 57 customer locations, and offers 38 products. Fig. 4 gives an overview of the relevant locations. Based on the data characteristics and data structure of the real-life problem, we create a large variety of test problems of different sizes. We coded our models and algorithms in CPLEX 12.1 and run the code on an Intel Xeon 5500 processor running at 2.66 GHz with 4 GB of memory. In Section 5.1, we test our solution approaches on differently sized problems, conducting extensive tests on the performance of our heuristics. In Section 5.2 we compare our comprehensive model with simplified formulations for different problem sizes.

5.1. Comparison of exact and heuristic solution approaches

Small and medium-sized problems can be solved well using an MIP solver. However, as problems grow, the MIP solver becomes less effective. Some very large problems cannot be solved using an MIP solver at all due to memory restrictions. Hence, heuristics like those introduced in Section 4 are needed. To analyze the performance of the various approaches, we create a number of test problems that are based on the data of the chemical company. The characteristics of the problems are summarized in Table 3. We generate 20 test problems per problem size, thus testing 240 different data sets overall.

We solve each problem with the MIP solver and with all heuristics, limiting the run time to one hour for all approaches. The results are summarized in Table 4. The lower bound we refer to in the table is the lower bound given by the MIP solver after one hour run time, i.e., after pruning the branch-and-bound tree. The MIP solver uses LP relaxations. Note that for problem 10, the MIP solver could not generate the problem due to memory limitations. We thus used the initial LP relaxation of the problem without pruning the branch-and-bound tree as a lower bound. Time refers to the convergence time (i.e., run time at which the best solution is found).

Table 3
Characteristics of random test problems for tests on heuristic performance.

Problem #	Number of variables	Thereof binary	PL	TF	CU	PR	T	B1	B3
1	73,232	4324	2	23	40	28	8	4	3
2	121,626	6264	2	29	50	34	8	4	3
3	186,340	8540	2	35	60	40	8	4	3
4	329,120	12,320	3	40	65	45	9	4	3
5	541,350	17,100	4	45	70	50	10	4	3
6	912,500	24,250	5	50	80	55	12	4	3
7	1,580,760	36,360	7	60	90	60	12	4	3
8	2,912,000	55,300	10	70	100	70	12	4	3
9	4,794,720	77,920	13	80	110	80	12	4	3
10	7,689,600	102,600	15	90	120	100	12	4	3

Table 4
Average gap of solutions to lower bound and best known solution and average convergence time. Results for best solution approach in bold.

Problem	MIP solver			D1			R1		
	Gap to LB [%]	Gap to best sol. [%]	Time [seconds]	Gap to LB [%]	Gap to best sol. [%]	Time [seconds]	Gap to LB [%]	Gap to best sol. [%]	Time [seconds]
1	**0.05**	**0.00**	3583	0.24	0.19	9	0.21	0.15	2213
2	**0.10**	**0.00**	3600	0.27	0.17	20	0.23	0.12	2244
3	**0.16**	**0.00**	3600	0.28	0.12	39	0.26	0.10	2169
4	**0.22**	**0.00**	3600	0.46	0.23	73	0.41	0.19	2260
5	**0.38**	**0.00**	3600	0.71	0.33	161	0.63	0.25	2349
6	**0.45**	**0.00**	3600	0.80	0.35	250	0.72	0.27	2405
7	1.60	0.90	3600	1.03	0.34	492	0.95	0.25	2516
8	1.72	0.97	3600	1.17	0.43	1190	1.10	0.35	2798
9	3.54	2.71	3600	1.31	0.50	2378	1.26	0.45	2891
10	n/a	n/a	n/a	1.33	0.50	3600	1.27	0.44	3591

	D2			R2		
	Gap to LB [%]	Gap to best sol. [%]	Time [seconds]	Gap to LB [%]	Gap to best sol. [%]	Time [seconds]
1	0.16	0.11	3062	0.17	0.12	3600
2	0.20	0.09	2885	0.20	0.10	3600
3	0.25	0.10	3068	0.27	0.12	3600
4	0.33	0.11	2911	0.35	0.12	3600
5	0.50	0.12	3274	0.53	0.14	3600
6	0.56	0.11	3271	0.59	0.14	3600
7	**0.70**	**0.01**	3477	0.72	0.03	3600
8	**0.75**	**0.01**	3484	0.77	0.02	3600
9	**0.81**	**0.01**	3600	0.81	0.01	3600
10	0.87	0.04	3600	**0.83**	**0.00**	3600

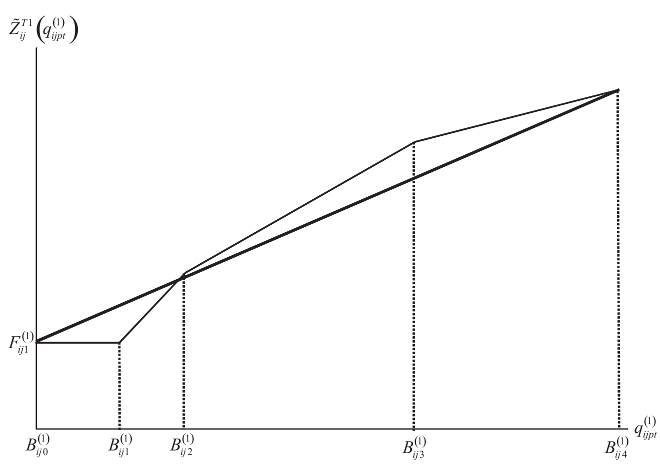

Fig. 5. Linearization technique for transport cost from plants to tank farms.

To evaluate the quality of the heuristics, we compare them in terms of solution quality and speed. In terms of solution quality, the MIP solver finds the best solutions of all approaches analyzed for small and medium-sized problems (Problems 1–6). For larger problems, heuristics D2 and R2 find lower cost solutions than the MIP solver. The largest problem, problem 10, cannot be solved using the MIP solver due to memory restrictions. The numerical analyses suggest that the MIP solver is the best choice for small and medium sized problems, but that the heuristics are a better choice for large problems.

We also compared the performances of four heuristics. The numerical results show that heuristics D2 and R2 achieve higher quality solutions than heuristics D1 and R1. This result is not surprising, because D2 and R2 model economy of scale effects more sophisticatedly than heuristics D1 and R1. Our numerical results suggest that the performance differences between both heuristics are small, but that D2 usually outperforms R2, though only marginally. However, heuristic D1 can be beneficial in situations where problem sizes are small and run times are critical. For small and medium-sized problems, D1 achieves reasonably good solutions at low run times.

To conclude our insights concerning the numerical analyses: for small and medium-sized problems, the MIP solver seems to be the best choice, but for larger problems D2 is superior.

Table 5
Comparison of different models for random test problems.

Prob. #	No. of variables	Thereof binary	PL	TF	CU	PR	T	B1	B3	Opt. gap (%)	Solution w/ M1	Δ M2 vs M1	Δ M3 vs M1
1	9220	970	2	10	20	15	4	4	3	0.0	388 (millon dollar)	+0.7 (millon dollar) (0.18%)	+2.1 (millon dollar) (0.54%)
2	51,888	3588	2	23	40	28	4	4	3	0.2	1579 (millon dollar)	+1.1 (millon dollar) (0.07%)	+3.4 (millon dollar) (0.22%)
3	140,420	7420	2	35	60	40	4	4	3	0.2	3464 (millon dollar)	+1.5 (millon dollar) (0.04%)	+5.0 (millon dollar) (0.14%)
4	218,720	9920	3	40	65	45	4	4	3	0.4	4300 (millon dollar)	+1.4 (millon dollar) (0.03%)	+7.1 (millon dollar) (0.16%)
5	321,030	12,780	4	45	70	50	4	4	3	0.4	4903 (millon dollar)	+1.9 (millon dollar) (0.04%)	+10.1 (millon dollar) (0.21%)

PL… plants, TF… tank farms, CU… customers, PR… products, T… transport schedules, B1… pieces of transport cost function from plant to tank farm, B3… pieces of warehousing cost function.

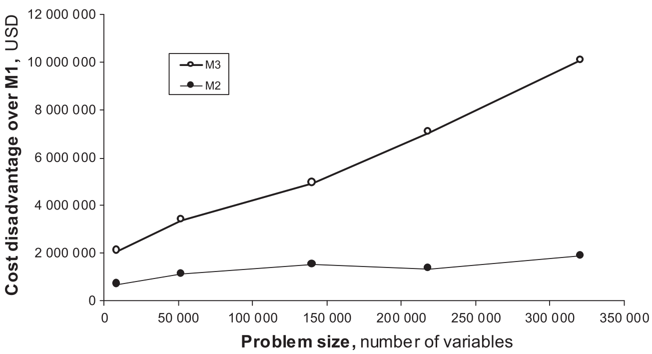

Fig. 6. Average cost disadvantage of models M2 and M3 over M1.

5.2. Effect of problem formulation on supply chain performance

In this section, we analyze the benefit of using a model that includes transport frequencies and non-linear cost functions as opposed to models that ignore some or all of these aspects. We refer to our comprehensive model as M1. We use two simple comparison models, M2 and M3. Model M2 contains cost functions that ignore quantity discounts (i.e., with the same fixed and variable cost factor irrespective of volume). Transport frequencies are optimized in this model. Model M3 ignores quantity discounts and does not optimize transport frequencies. In M3, only one transport frequency is possible. We choose the transport frequency most heavily used in the solution obtained by model M1 when testing M3. Both M2 and M3 do not consider quantity discounts. It is thus necessary to apply some linear approximation to the cost functions. We approximated the original, piecewise linear cost functions $Z(q)$ using the chord between the points $(0;F_1)$ and $(B_N;Z(B_N))$, i.e., averaging the variable cost portion over the complete domain. See Geoffrion (1977) for a similar technique. We tested several other naïve linearization methods, but found this one performed best. Fig. 5 provides an example of this technique for $\widetilde{Z}^{T1}_{ijt}$, the transport cost from plants to tank locations. We hence reformulate the model using these cost functions instead of the complex non-linear versions introduced in Section 3.

We test five data sets each for five different problem sizes, ranging from small to large. Our objective is analyzing the effect of the model complexity on performance. To isolate this effect from other effects, we solve all problems with a standard branch-and-bound code. Due to the size of the problems, solving to optimality is not feasible within a reasonable time for all problem sizes. Table 5 gives an overview of the problem sizes analyzed, the optimality gap we deemed acceptable for our calculations, the cost of the best solution obtained with M1, and the absolute and relative cost disadvantage obtained with M2 and M3.

We evaluate the solutions obtained by model M2 and M3 using the original cost structures of model M1. Our experiments show that model M1 leads to substantial cost savings over the simplified models. For example, for a problem with 4 plants, 45 tank farms, 70 customers, 50 products, and 4 transport schedules, M1 has cost advantages amounting to up to 10 million USD per year over M3. Cost advantages over model M2 still amount to up to 2 million USD per year. Fig. 6 gives an overview of the average results.

6. Conclusion

In this paper, we have developed a supply chain design model exhibiting economies of scale in transport and warehousing that explicitly takes transport frequencies into consideration. As our tests on real data show, failing to consider economies of scale and transport frequencies leads to significantly higher costs. Due to their mixed-integer nature, realistically sized problems are very time consuming to solve optimally. For very large problem instances, they are even impossible to solve with standard optimization software. We have developed two solution heuristics with deterministic and stochastic variants based on iterative linearization techniques, exploiting the efficient optimization methods available for linear problems. While standard branch-and-bound codes perform well on small to medium-sized problems, like the problem that motivated our research, our heuristics outperform them in terms of solution quality for large problems within the set time limits. Regarding solution time, our heuristics are superior even for small problems.

While we assumed an uncapacitated supply chain, there can exist capacity limits on production output, vessel size, and tank size in real-life situations. The capacitated case can be easily implemented. The model formulation requires no adaptation at all. The upper bound of the respective cost functions is simply set to the capacity limit instead of a ‘‘Big M.’’ The heuristics can also deal with capacitated problems with some alterations to ensure that the problem does not become infeasible by dropping tank farms from the solution.

Different supply chain structures with elements such as cross-docking or the possibility of direct deliveries from plants to customers can also be included in the model. Both require additional constraints. In the cross-docking case, we need to ensure that the transport schedules from the plants to the cross-docks are the same as those from the cross-docks to the customers. For direct deliveries, a constraint is needed to exclude shipping schedules that do not comply with customer demands.

In some situations, volume discounts also apply to the second transport step from tank farm to customer. For example, this may be the case for problems where this transport step is performed using trucks, where discounts already play a role for smaller quantities than in sea transport. In this case, the cost function in the model formulation is replaced with a piecewise linear function. In general, many different kinds of cost functions can be used in our model since concavity is not required. This includes staircase functions, functions exhibiting diseconomies of scale, etc.

In our model, transport frequencies for shipments from tank farms to customers are fixed to the same schedule for all customers. However, the model and the proposed solution approaches can be adjusted to scenarios with different shipping schedules for different customers by replacing the parameter $O^{(2)}$ with $O^{(2)}_k$, thus feeding a different schedule into the model for each customer. Additional constraints are needed to ensure that the transport schedules from plants to tank farms are in sync with those from tank farms to customers. Schedules are in sync when transport frequencies from tank farms to customers are a multiple of transport frequencies from plants to tank farms (e.g., delivery from tank farms to customers every week; delivery from plants to tank farms every four weeks). Otherwise, situations may arise where an insufficient quantity of chemicals is available at a tank farm when a shipping to customers from this tank farm is due. Taking our thinking one step further, it is also possible to optimize transport frequencies for the second transport step: from tank farms to customers. In this case, the index t is added to the decision variables for this transport step. Again, constraints to ensure the synchronization of schedules between transport from plants to tank farms and transport from tank farms to customers are necessary.

References

Balakrishnan, A., Graves, S.C., 1989. A composite algorithm for a concave-cost network flow problem. Networks 19, 175–202.

BASF, 2007. Company website – News and Media Relations. <http://corporate.basf.com/de/presse/?id=WGK-HBY8Wbcp*xP> (accessed 15.03.08).

Baumol, W.J., Wolfe, P., 1958. A warehouse-location problem. Operations Research 6, 252–263.
Baxter, J., 1984. Depot location: A technique for the avoidance of local optima. European Journal of Operational Research 18, 208–214.
Bayer, 2007. Bayer will Kapazität der geplanten TDI-Produktion in Shanghai auf 300.000 Jahrestonnen erweitern (February 6, 2007). <http://www.investor.bayer.de/de/news/investor-news/investor-news/showNewsItem/738/1170745507/5dc1abb86f/> (accessed 15.03.08).
Broek, J., Schütz, P., Stougie, L., Tomasgard, A., 2006. Location of slaughterhouses under economies of scale. European Journal of Operational Research 175, 740–750.
Correia, I., Captivo, M.E., 2003. A Lagrangian heuristic for a modular capacitated location problem. Annals of Operations Research 122, 141–161.
Drysdale, J.K., Sandiford, P.J., 1969. Heuristic warehouse location – A case history using a new method. INFOR 7, 45–61.
Efroymson, M.A., Ray, T.L., 1966. A branch-bound algorithm for plant location. Operations Research 14, 361–368.
Feldman, E., Lehrer, F.A., Ray, T.L., 1966. Warehouse location under continuous economies of scale. Management Science 12, 670–684.
Fleischmann, B., 1993. Designing distribution systems with transport economies of scale. European Journal of Operational Research 70, 31–42.
Geoffrion, A.M., 1977. Objective function approximations in mathematical programming. Mathematical Programming 13, 23–37.
Gümüs, M., Bookbinder, J.H., 2004. Cross-docking and its implications in location-distribution systems. Journal of Business Logistics 25, 199–228.
Holmberg, K. 1984. Capacitated facility location with staircase costs. Research Report LiTH-MAT-R-1984-26, Department of Mathematics, Linkoping Institute of Technology, Sweden.
Holmberg, K., 1994. Solving the staircase cost facility location problem with decomposition and piecewise linearization. European Journal of Operational Research 75, 41–61.
Holmberg, K., Ling, J., 1997. A Lagrangean heuristic for the facility location problem with staircase costs. European Journal of Operational Research 97, 63–74.
Kelly, D.L., Khumawala, B.M., 1982. Capacitated warehouse location with concave costs. The Journal of the Operational Research Society 33, 817–826.
Khumawala, B.M., Kelly, D.L., 1974. Warehouse location with concave costs. INFOR 12, 55–65.
Kim, D., Pardalos, P.M., 2000a. Dynamic slope scaling and trust interval techniques for solving concave piecewise linear network flow problems. Networks 35, 216–222.
Kim, D., Pardalos, P.M., 2000b. A dynamic domain contraction algorithm for nonconvex piecewise linear network flow problems. Journal of Global Optimization 17, 225–234.
Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P., 1983. Optimization by simulated annealing. Science 220, 671–680.
Klincewicz, J.G., 1985. A large-scale distribution and location model. AT&T Technical Journal 64, 1705–1730.
Klincewicz, J.G., 1990. Solving a freight transport problem using facility location techniques. Operations Research 38, 99–109.
Klincewicz, J.G., 2002. Enumeration and search procedures for a hub location problem with economies of scale. Annals of Operations Research 110, 107–122.
Klincewicz, J.G., Luss, H., Yu, C.S., 1988. A large-scale multilocation capacity planning model. European Journal of Operational Research 34, 178–190.
Lapierre, S.D., Ruiz, A.B., Soriano, P., 2004. Designing distribution networks: Formulations and solution heuristics. Transportation Science 38, 174–187.
Lin, J.R., Nozick, L.K., Turnquist, M.A., 2006. Strategic design of distribution systems with economies of scale in transportation. Annals of Operations Research 144, 161–180.
O'Kelley, M.E., Bryan, D.L., 1998. Hub location with flow economies of scale. Transportation Research Part B 8, 605–616.
Paraschis, I.N., 1989. Optimale Gestaltung von Mehrprodukt-Distributionssystemen: Modelle – Methoden – Anwendungen. Physica-Verlag, Heidelberg.
Romeijin, H.E., Shu, J., Teo, C.P., 2007. Designing two-echelon supply networks. European Journal of Operational Research 178, 449–462.
SABIC, 2006. Annual Report 2006.
SABIC, 2007. Company website – News and Media Relations. <http://www.sabic.com/corporate/en/newsandmediarelations/news/default.aspx> (accessed 15.03.08).
Soland, R.M., 1974. Optimal facility location with concave costs. Operations Research 22, 373–382.
Spielberg, K., 1969a. An algorithm for the simple plant location problem with some side conditions. Operations Research 17, 85–111.
Spielberg, K., 1969b. Plant location with generalized search origin. Management Science 16, 165–178.
Syam, S.S., 2002. A model and methodologies for the location problem with logistical components. Computers and Operations Research 29, 1173–1193.
Whitaker, R.A., 1985. Some add-drop and drop-add interchange heuristics for non-linear warehouse location. Journal of the Operational Research Society 36, 61–70.
Zangwill, W.I., 1968. Minimum concave cost flows in certain networks. Management Science 14, 429–450.